



HAIChart: Human and AI Paired Visualization System

Yupeng Xie
HKUST (GZ)

yxie740@connect.hkust-gz.edu.cn

Yuyu Luo*

HKUST (GZ) / HKUST
yuyuluo@hkust-gz.edu.cn

Guoliang Li

Tsinghua University
liguoliang@tsinghua.edu.cn

Nan Tang

HKUST (GZ) / HKUST
nantang@hkust-gz.edu.cn

ABSTRACT

The growing importance of data visualization in business intelligence and data science emphasizes the need for tools that can efficiently generate meaningful visualizations from large datasets. Existing tools fall into two main categories: human-powered tools (e.g., Tableau and PowerBI), which require intensive expert involvement, and AI-powered automated tools (e.g., Draco and Table2Charts), which often fall short of *guessing* specific user needs.

In this paper, we aim to achieve the best of both worlds. Our key idea is to initially auto-generate a set of high-quality visualizations to minimize manual effort, then refine this process iteratively with user feedback to more closely align with their needs. To this end, we present HAICART, a reinforcement learning-based framework designed to iteratively recommend good visualizations for a given dataset by incorporating user feedback. Specifically, we propose a Monte Carlo Graph Search-based visualization generation algorithm paired with a composite reward function to efficiently explore the visualization space and automatically generate good visualizations. We devise a visualization hints mechanism to actively incorporate user feedback, thus progressively refining the visualization generation module. We further prove that the top-*k* visualization hints selection problem is NP-hard and design an efficient algorithm. We conduct both quantitative evaluations and user studies, showing that HAICART significantly outperforms state-of-the-art human-powered tools (21% better at Recall and 1.8× faster) and AI-powered automatic tools (25.1% and 14.9% better in terms of Hit@3 and R10@30, respectively).

PVLDB Reference Format:

Yupeng Xie, Yuyu Luo, Guoliang Li, and Nan Tang. HAIChart: Human and AI Paired Visualization System. PVLDB, 17(11): 3178 - 3191, 2024.
doi:10.14778/3681954.3681992

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/HKUSTDial/HAIChart>.

1 INTRODUCTION

Data visualization is an effective means to uncover underlying insights [41, 46, 61, 68]. With the increasing importance of data visualization, *how to help users effectively and easily create visualizations from massive datasets* has attracted extensive attention from academia [22, 24, 28, 29, 47, 54, 65] and industry (e.g.,

*Yuyu Luo is the corresponding author.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 17, No. 11 ISSN 2150-8097.
doi:10.14778/3681954.3681992

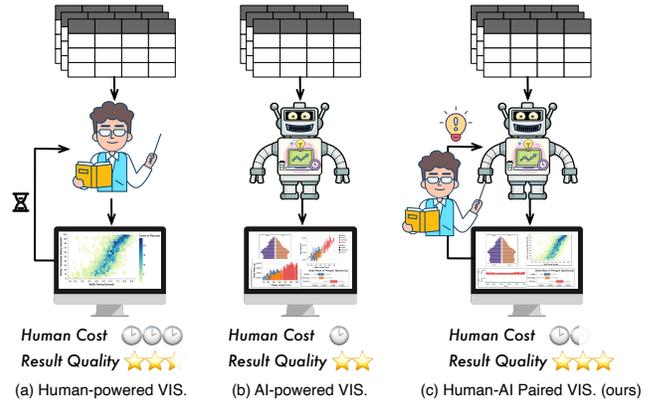


Figure 1: Three cases of visualization approaches

Tableau [2]). Existing visualization tools can be broadly categorized into two groups based on the user effort required for creating visualizations, i.e., *human-powered* and *AI-powered* tools.

Human-powered Visualization tools, including Voyager2 [63], DeVIL [66], and Lary [44], enable users to create their desired visualizations by manually specifying the data attributes, transformations, and visual encodings, as shown in Figure 1(a). This process is time-consuming and error-prone because it requires a deep understanding of the datasets and analytical tasks. Users often need to attempt multiple iterations to achieve the final visualization. Consequently, interactive visualization demands *domain-* and *data-specific* expertise, which can significantly hinder novice users from effectively engaging in visual analysis.

In response to the above challenges, recent research tries to *automate* the visualization process empowered by artificial intelligence.

AI-powered Automatic Visualization aims to enumerate and recommend the (top-*k*) best visualizations for a given dataset based on predefined constraints [13, 26, 31, 36, 38] or learning-based recommendation algorithms [39, 69]. For example, DeepEye [31] recommends visualizations based on visualization rules and a ranking model, while Table2Charts [69] generates top-*k* visualizations using a table-to-sequence generation model. Although these tools provide a valuable starting point for data analysis by suggesting “good” visualizations, they risk misleading users with indiscriminate recommendations. The key issue is that these tools typically rely on static logical rules or deep learning models without adequately capturing user intent or feedback, hindering the tailoring of recommended visualizations to better suit users’ needs.

Human and AI Paired Visualization. To mitigate the limitations of the above methods and strike a balance between human-powered and AI-powered visualization, we present *Human and AI Paired*

Table 1: Comparison of HAICHART with existing works (Dataset: D , User Operations: U , Visualizations: V)

Types	Systems	Input	Output	Multi-round?	Learned?	Recommendation Perspectives		
						Domain Knowl.	User Preferences	Data Features
Human-powered	Voyager2 [63]	D, U	V	\times	\times	\checkmark	\times	\times
	DeVIL [66]	D, U	V	\times	\times	\checkmark	\checkmark	\times
	Lary [44]	D, U	V	\times	\times	\checkmark	\checkmark	\times
AI-powered	DeepEye [31]	D	V	\times	\checkmark	\checkmark	\times	\checkmark
	VizML [19]	D	V	\times	\checkmark	\times	\times	\checkmark
	Draco-Learn [38]	D	V	\times	\checkmark	\checkmark	\checkmark	\checkmark
	Data2Vis [15]	D	V	\times	\checkmark	\times	\times	\checkmark
	KG4VIS [26]	$D, \text{Knowl. Graph}$	V	\times	\checkmark	\checkmark	\times	\times
	Table2Charts [69]	D	V	\times	\checkmark	\times	\times	\checkmark
	VizGRank [16]	D	V	\times	\times	\checkmark	\times	\checkmark
	PVisRec [39]	D	V	\times	\checkmark	\times	\checkmark	\checkmark
	LLM4Vis [59]	D	V	\times	\checkmark	\times	\times	\checkmark
	PI2 [10]	$D, \text{SQL Queries}$	$V, \text{Interface}$	\checkmark	\times	\checkmark	\checkmark	\times
Human-AI Paired	HAICHART	$D, \text{User Feedback}$	V, Hints	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark

Visualization, as shown in Figure 1(c). Our key idea is to *initially* recommend a set of *high-quality* visualizations and then continue to refine these recommended visualizations by providing users with *visualization hints* for further tuning. This approach not only simplifies the process for novice users seeking to discover effective visualizations but also empowers users to actively participate in and guide the AI-powered visualization process through the human-centered feedback, as shown in Table 1.

Note that another natural way humans interact with AI for visualizations is through conversations powered by Large Language Models (LLMs), or briefly LLM4VIS [11, 14, 18, 21, 56, 59]. However, the key limitation of LLM4VIS is that even with user natural language hints, it is hard to precisely modify visualizations outputted by LLMs using natural language; that is, easy-to-use but hard-to-calibrate. Please refer to Section 7 for an empirical comparison with LLM4VIS powered by GPT-4 and Section 8 for more discussions.

Challenges. There are three challenges in our problem. (C1) How can we effectively and efficiently explore the visualization search space to recommend visualizations that are both high-quality and relevant? (C2) How can we evaluate the “goodness” of a generated visualization comprehensively? (C3) How can we understand and effectively integrate user feedback to guide the system toward visualizations that align with users’ requirements?

Our Proposal: HAICHART. To address these challenges, we propose HAICHART to automatically and progressively recommend a series of visualizations that adapt to evolving user preferences over time. This process is characterized by the system’s ability to learn from user interactions and feedback, iteratively refining the recommended visualizations. *The objective is to converge towards a final set of visualizations that closely align with the user’s preferences, within the constraints of the available data and visualization search space.* This iterative optimization is essentially a series of decisions based on user intent, similar to the key idea in reinforcement learning.

We adopt the reinforcement learning framework to implement HAICHART. We treat a visualization as a sequence of operations, such as chart types and x/y -axes configurations. Thus, the “state” is the current sequence of *incomplete* visualization, the “action” refers to selecting the subsequent visualization operation within that state, and the “environment” is the visualization system itself.

The environment evaluates the quality of the visualizations and provides rewards to guide the “agent” responsible for generating and recommending visualizations.

We utilize the Monte Carlo Graph Search (MCGS) algorithm to build this agent, enabling it to efficiently navigate the extensive search space to discover optimal visualizations (addressing C1). To more accurately evaluate the visualizations, we design a composite reward function that takes into account the data features, visualization domain knowledge, and user preferences to ensure the quality of the visualizations (addressing C2). Moreover, we devise a visualization hints selection algorithm for recommending useful hints (e.g., “explore why flights are delayed”) to assist the user in data exploration. This approach can integrate user feedback for better visualization results (addressing C3).

Contributions. We make the following notable contributions.

- (1) Problem Statement. We formally define the problem of human and AI paired visualization (Section 2).
- (2) HAICHART. We propose a reinforcement learning-based system, HAICHART, to integrate human insights and AI capabilities for better visualizations. HAICHART initially generates a set of promising visualizations for users and iteratively improves the alignment of recommended visualizations with user preferences by selecting a set of visualization hints for interaction (Sections 3 & 4).
- (3) Learning-to-Rate Visualization. We design a composite reward function to rate the quality of visualizations, which is used to accurately guide the visualization generation process (Section 5).
- (4) Visualization Hints Selection. We design visualization hints as the proxy to collect user feedback, prove the selection of the top- k visualization hints is an NP-hard problem and design an efficient algorithm to select the top- k hints (Section 6).
- (5) Experiments. We conduct both quantitative evaluations and user studies to demonstrate that HAICHART outperforms human-powered tools in terms of both accuracy (21% better at Recall) and efficiency in creating visualizations (1.8× faster) and is better than AI-powered systems in effectiveness by 25.1% and 14.9% in terms of Hit@3 and R10@30, respectively (Section 7).

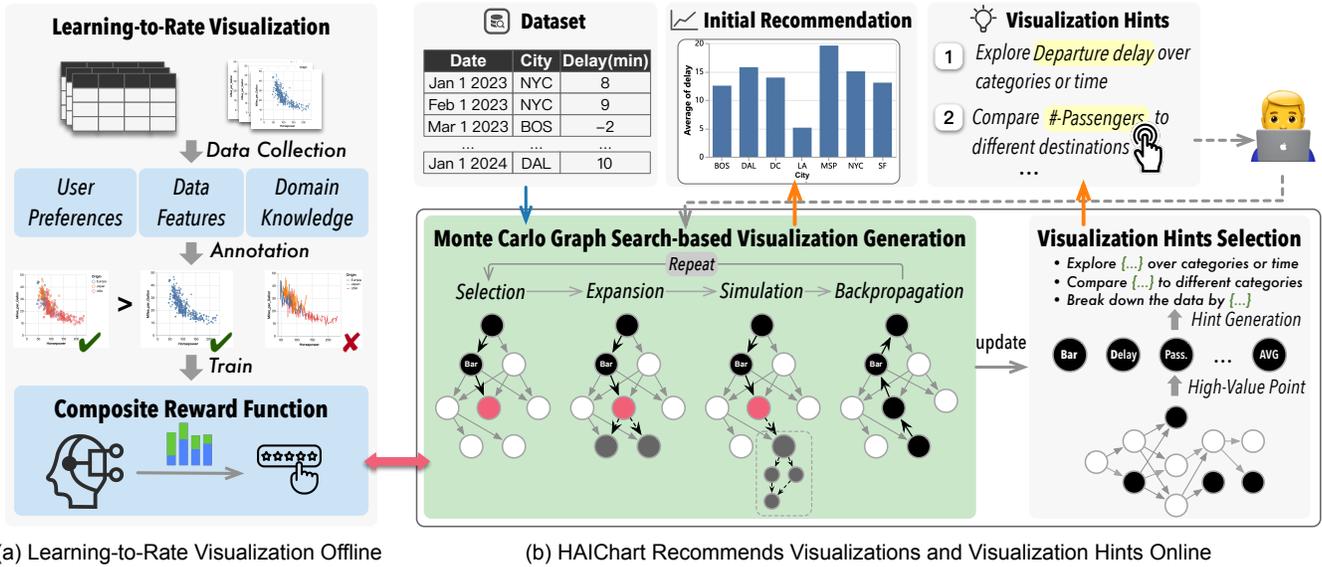


Figure 2: An overview of HAICART

2 PROBLEM AND SOLUTION OVERVIEW

2.1 Problem

Given a relational table D , let \mathcal{V} be all valid visualizations, and \mathbf{V}^+ be a set of “good” visualizations, both being associated with D . In practice, this set of visualizations \mathbf{V}^+ is often manually selected by the users, such as the visualizations in a dashboard.

Visualization Selection. The problem of *visualization selection* is to select a set \mathbf{V} of visualizations from all valid visualizations \mathcal{V} , which maximizes $\mathbf{V} \cap \mathbf{V}^+$ (i.e., maximizes the number of good visualizations being selected) and minimizes $\mathbf{V} \setminus \mathbf{V}^+$ (i.e., minimizes the number of bad visualizations being selected).

When being operated manually, visualization selection is time-consuming, and its effectiveness depends on the user’s expertise in data visualization and domain knowledge of specific datasets.

To reduce human effort, visualization recommendation algorithms have been studied. These algorithms are mostly powered by machine learning algorithms. However, these recommendation algorithms are mainly one round.

One-Round Visualization Recommendation. The problem of *one-round visualization recommendation* is to automatically suggest a set \mathbf{V}^S of visualizations and ask the user to select a subset $\mathbf{V} \subseteq \mathbf{V}^S$.

In practice, a one-round visualization recommendation may not be sufficient to support the user’s diversified intent, e.g., selecting all charts for a dashboard. Therefore, we propose to study a new problem, which does multiple rounds of interactions with users to collectively select all required good visualizations.

Multi-Round Visualization Recommendation. The problem of *multi-round visualization recommendation* is to recommend n rounds of visualizations to users. In each round, we can suggest \mathbf{V}_i^S visualizations (optionally with some natural language hints), and the user can select a subset $\mathbf{V}_i \subseteq \mathbf{V}_i^S$ visualizations in the i -th round. After n rounds, the user will select $\mathbf{V} = \bigcup_{i=1}^n \mathbf{V}_i$ visualizations.

2.2 Solution Overview

HAICART Overview. As shown in Figure 2, HAICART consists of two components: an offline part for learning-to-rate visualizations and an online part for automatically recommending good visualizations and providing visualization hints for user interaction.

Offline: Learning-to-Rate Visualization. As shown in Figure 2(a), the offline part is responsible for understanding and evaluating visualization quality. To achieve this, we incorporate visualization rule of thumb, data features, and user preferences to learn a composite reward function.

Online: Multi-Round Visualizations and Hints Recommendation. As shown in Figure 2(b), HAICART leverages an *agent* powered by Monte Carlo Graph Search to recommend a set of promising visualizations along with visualization hints to guide further exploratory actions. Specifically, this agent initially traverses the visualization search space effectively and efficiently. It generates possible visualizations, estimates their quality based on the well-trained composite reward function, and finally returns a ranked list of high-quality visualizations. In addition, the system computes visualization hints, derived from the current states of the Monte Carlo Graph. The user can browse the recommended visualizations and hints. Once the user selects a hint, the system will further explore the visualization space and recommend visualizations guided by the selected hint.

3 REINFORCEMENT LEARNING FOR MULTI-ROUND VISUALIZATION RECOMMENDATION

3.1 Preliminary

Inspired by previous works [31, 36, 38], we adopt a visualization query to represent all possible visualizations discussed in this paper.



Figure 3: A visualization query

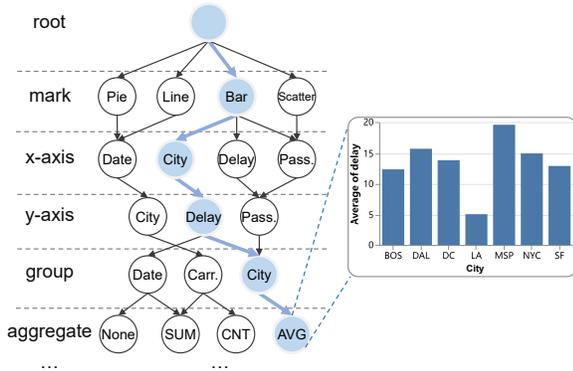


Figure 4: A visualization query graph

Visualization Query. Given a relational table D with the scheme $R(A_1, \dots, A_m)$, a visualization query Q is composed of a sequence of visualization operations (*i.e.*, actions) \mathcal{A} (mark, encoding, transformation). Specifically:

- mark represents the type of visualization charts including bar, line, pie, and scatter.
- encoding denotes the encoding of data attributes including x/y -axes, aggregate functions, and color schemes.
- transformation defines data transformations: filtering, binning, grouping, sorting, and top- k operations.

Applying the visualization query Q to the table D results in the corresponding visualization V , *i.e.*, $V = Q(D)$. Therefore, any visualization V for a table D can be represented by a combination of data attributes and cell values in D , along with visualization operations (*i.e.*, actions) in \mathcal{A} .

Example 1: [Visualization Query] Figure 3 shows an illustrative query (Q) applied to the flight delays dataset (D). The query visualizes a bar chart with the x -axis encoding cities (*i.e.*, City column) and the y -axis showing the average delay (*i.e.*, $\text{AVG}(\text{Delay})$). The bar chart ($Q(D)$) shows the trend in average delays, showing that LA (*i.e.*, Los Angeles) has the shortest average delays, while MSP (*i.e.*, Minneapolis) has the longest average delays. \square

Undoubtedly, the search space for visualizations is very large, growing exponentially with the increase in number of columns and combinations of visualization operations.

To effectively navigate this huge search space, we introduce the concept of *visualization query graph* to represent all possible visualizations for a given dataset, defined as follows:

Definition 3.1: [Visualization Query Graph] Given a table D , the visualization query graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ is defined as a directed acyclic

graph. Specifically, each node $v \in \mathcal{V}$ represents a visualization operation, and each directed edge $e \in \mathcal{E}$ indicates a transition from node v_i to v_j . The weight of an edge (v_i, v_j) , denoted by w_{ij} , indicates the effectiveness of transitioning from operation v_i to v_j . A path from the root node to an end node within this graph represents a sequence of visualization query operations, which together form a candidate visualization query. \square

Example 2: [Visualization Query Graph] As illustrated in Figure 4, each layer of the graph corresponds to the possible values that a visualization action can take. For example, the chart type (*mark*) could be *bar*, *line*, *scatter*, *pie*, etc. Taking the blue path as an example, it represents a visualization query – “**mark Bar encoding x City y AVG(Delay) transform group City**”. After executing this query over the dataset, it will create a bar chart. \square

3.2 HAICHAART Details

Given a dataset for visualization, users typically select visualization types, data columns, and transformation operations based on their preliminary understanding. If the results do not meet their analysis requirements, they adjust their operations to optimize the output. This visualization process is a sequence of decision-making based on user intent, similar to Reinforcement Learning (RL), where feedback from the visualization system (*i.e.*, environment) helps optimize strategies to achieve satisfactory outcomes.

Therefore, we model the problem of visualization generation and recommendation as a Markov Decision Process (MDP) [6] and implement it using an RL framework.

The key RL components of HAICHAART are as follows:

State. To apply the RL framework to our problem, it is crucial to accurately define the state, which serves as input to the agent and aids in making decisions during visualization generation. We transform the visualization generation problem into a process that starts from an initial state and proceeds through a series of visualization operation/action decisions to reach a target state. Each state corresponds to a visualization query consisting of multiple visualization clauses, as described in Section 3.1.

Specifically, visualization queries can be classified into two types:

(1) *Partial Visualization Queries.* These queries need to be further extended to form a complete query and reflect historical decision paths in Visualization Query Graph. For example, an incomplete query might have specified the type of chart (*e.g.*, bar chart) but not yet selected a specific data field for the x - or y -axis.

(2) *Complete Visualization Queries.* These queries do not require additional extensions and can be evaluated for visualization quality by our reward function. After that, the system updates the nodes and edges of the graph based on the evaluation results.

Action. An action is an operation that an agent can perform based on its current state, determining the next step in a visualization query sequence. As discussed in Section 3.1, the visualization query can be categorized into three main types of actions: mark, encoding, and transformation. Therefore, for a given dataset D , the available action space \mathcal{A} is fixed.

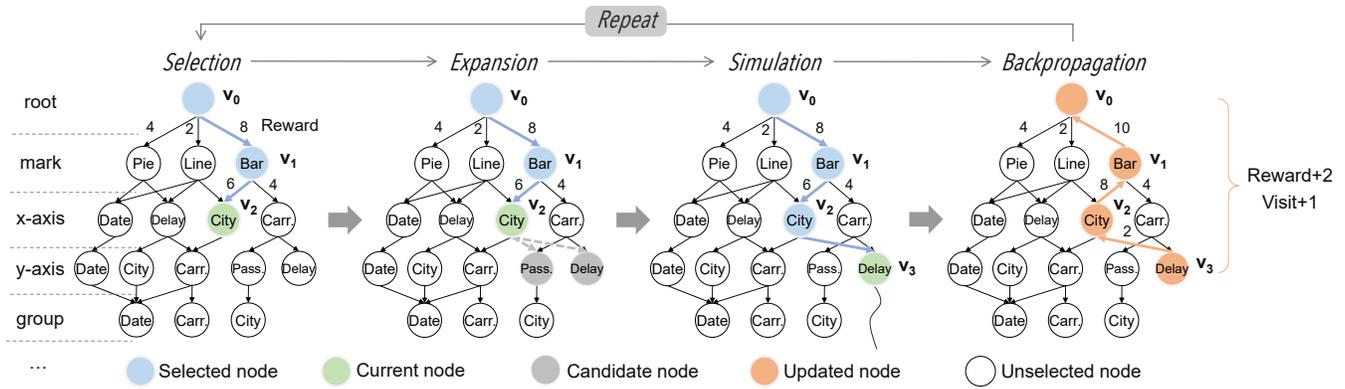


Figure 5: An example of MCGS-based visualization generation

Algorithm 1: MCGS-based Visualization Generation

```

Input: Dataset  $D$ ;
Output: Top- $k$  Visualization Results  $R = \{V_1, \dots, V_k\}$ ;
1  $node \leftarrow \text{Initialize}(D)$ ;  $\mathcal{G} \leftarrow \{node\}$ ;  $R \leftarrow \{\}$ ;
2 for each iteration do
3   Set a visualization query  $Q = ""$ ;
4    $node = root$ ;
5   while  $Q$  is invalid do
6     if  $node$  is fully expanded then
7        $node \leftarrow \text{select}(\mathcal{G}, node)$ ;
8       Add  $node$  to  $Q$ ;
9       if  $Q$  is valid then
10        break;
11    else
12       $node \leftarrow \text{expand}(\mathcal{G}, node)$ ;
13       $Q \leftarrow \text{simulation}(\mathcal{G}, node)$ ;
14   $score \leftarrow \text{reward}(Q)$ ;
15   $\text{BackPropagation}(score, Q, \mathcal{G})$ ;
16   $V \leftarrow Q(D)$ ; // Corresponding visualization
17  add  $V$  to  $R$ ;
18 return  $R$ ;

```

Agent. The agent combines the Monte Carlo Graph Search (MCGS) algorithm and the Upper Confidence Bound (UCB) [4] algorithm to generate visualizations and make decisions. The MCGS performs well in navigating complex and uncertain search spaces, while the UCB algorithm effectively balances the trade-off between exploiting known optimal solutions and exploring new possibilities, thereby optimizing the agent’s decision-making process (see Section 4).

Reward. The reward function evaluates the quality of the visualizations created by the MCGS algorithm and updates the visualization query graph accordingly. This guides the agent to generate higher quality visualizations. The reward function considers three key factors: data features, visualization domain knowledge, and user preferences to accurately assess visualization quality (see Section 5).

Environment. The environment in HAICART first uses a pruning algorithm to efficiently generate valid visualization queries during

the generation process. Second, it evaluates the quality of valid queries and returns a reward to guide the training process.

4 VISUALIZATION GENERATION

4.1 Monte Carlo Graph Search-based Visualization Generation

In Section 3.1, we introduced the visualization query graph, where each visualization query is represented by a path within the graph. All possible paths within the query graph constitute the search space for visualizations. We can use search algorithms to explore the graph and find high-quality visualizations. However, as the dataset size increases, the search space expands exponentially, making it challenging to efficiently find good visualization results.

Traditional graph search methods like Depth-First Search (DFS) and Breadth-First Search (BFS) perform well in structured search spaces but often fall short in large and complex spaces because they cannot dynamically adjust the search strategy based on the information already found. To address these limitations, we propose a Monte Carlo Graph Search-based algorithm to navigate the vast search space efficiently.

Key Idea. Our key idea is to dynamically accumulate and utilize shared information from nodes throughout the search process. Initially, shared information is limited, but as the search progresses, the accumulated information from explored nodes becomes increasingly significant, providing greater support for decision-making. This allows the algorithm to gradually reduce its reliance on random simulations, leveraging accumulated knowledge to guide the search and enhance efficiency [23].

We first overview our Monte Carlo Graph Search-based visualization generation algorithm. As shown in Algorithm 1, the algorithm begins by taking a dataset D as input, aiming to generate the top- k visualization results R . In the initial phase, the algorithm initializes a root node based on the dataset D , which serves as the starting node of the graph \mathcal{G} (Line 1).

In each iteration, the algorithm starts by initializing a visualization query Q (Line 3) and then starts exploring from the root node. During the node selection phase, each choice is added to the visualization query Q (Line 7-8) until an unexpanded node is encountered,

followed by an expansion of this node (Line 12). When the visualization query Q is valid, the algorithm utilizes a reward function to calculate the score of the query (Line 14), then performs back-propagation to update the information on the graph \mathcal{G} (Line 15). Finally, the query is transformed into the corresponding visualization results and stored in R (Lines 16-17). After k visualizations are selected, the algorithm returns this set R .

Next, we will detail the procedure of Algorithm 1 through Figure 5. The algorithm mainly includes the following four steps:

Selection. In the selection phase (Line 7 in Algo. 1), the algorithm starts from the root node and recursively selects the optimal child nodes until it reaches a node that has not yet been fully expanded. As shown in Figure 5, the blue nodes represent the nodes selected in this phase, while the green nodes represent the currently selected nodes that have not been fully expanded. To effectively utilize the feedback information provided by the reward function and balance exploration and exploitation, we adopt the Upper Confidence Bound (UCB) algorithm. The key idea of this algorithm is to select the child node with the highest confidence upper bound at each iteration. The specific selection strategy of the UCB algorithm can be represented by the following formula:

$$UCB = \underbrace{\bar{X}_i}_{\text{exploitation}} + \underbrace{c\sqrt{2\ln n/n_i}}_{\text{exploration}} \quad (1)$$

where \bar{X}_i is the average reward of child node i , n is the number of visits to the current node, n_i is the number of visits to child node i , and c is a constant used to balance exploration and exploitation.

Expansion. During the expansion phase (Line 12 in Algo. 1), the algorithm selects the next valid action based on the current state (*i.e.*, an incomplete visualization query). Specifically, it removes those low-quality visualizations that are either syntactically incorrect or violate visualization rules (see Section 4.2). For example, when processing node v_2 (in Figure 5), the algorithm identifies two high-benefit candidates and randomly selects one for further expansion.

Simulation. During the simulation phase (Line 13 in Algo. 1), the algorithm starts at the current node (*e.g.*, node v_3 shown in Figure 5) and performs simulation actions. First, it checks whether the current visualization query is valid. If the query is invalid, the algorithm randomly explores the next node according to the visualization grammar rule until a valid query is constructed. Once a valid query is formed, the learning-to-rate visualization part assigns a reward.

Backpropagation. This phase (Line 15 in Algo. 1) propagates the simulation results through the visualization query graph. After the simulation phase, as shown in Figure 5, the score obtained by node v_3 is propagated along its path. Each node and edge along the path updates its reward value and visit count based on the simulation outcomes. Our learning-to-rate visualization module evaluates the visualization quality and updates the node information along the search path, guiding the algorithm to make more precise decisions in subsequent explorations by feeding the results of each simulation back into the visualization query graph.

Terminal Condition. The above steps are repeated until the maximum number of iterations is reached.

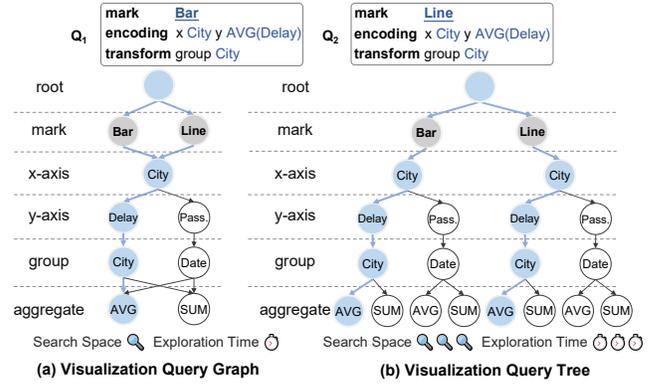


Figure 6: Comparison of tree and graph structures

Comparison of MCGS and MCTS. The Monte Carlo method solves complex problems through random sampling but requires simulation steps for each new node, impacting performance based on node count [42]. Traditional MCTS methods rely on extensive, random, and time-consuming simulations to determine the next step [49]. Unlike traditional MCTS, our MCGS algorithm explores a graph structure instead of a tree structure, allowing for more effective information sharing between nodes and reducing the number of nodes, thereby increasing search efficiency [57].

As shown in Figure 6, visualization queries Q_1 and Q_2 differ only in chart type; other clauses are identical. In a tree structure (Figure 6(b)), this information cannot be shared due to the strict parent-child hierarchy, leading to redundancy. In contrast, the graph structure (Figure 6(a)) allows node information sharing across different queries, reducing redundancy and improving search efficiency.

More concretely, we compare the number of nodes between MCGS and MCTS. Consider 4 chart types, a dataset with m columns, encoding that considers both the x - and y -axes, and data transformations that include grouping and 4 aggregate operations. In MCTS, each path is independent, resulting in a total node count of $N_{\text{tree}} = 4 \times m^3 \times 4 = 16m^3$. In MCGS, using the graph structure to share nodes, the total node count is $N_{\text{graph}} = 4 + 3 \times m + 4 = 3m + 8$. The reduction factor is $R = \frac{N_{\text{tree}}}{N_{\text{graph}}} = \frac{16m^3}{3m+8}$. As m increases, the benefit of using MCGS becomes more significant, as it reduces the number of nodes compared to MCTS.

In summary, by sharing nodes and reducing redundant calculations, MCGS significantly reduces the number of nodes and simulations, enabling more efficient search [12, 43, 48].

4.2 Optimization Techniques

Rule-based Pruning. In the visualization generation process, selecting and expanding nodes are crucial. However, the traditional UCB algorithm may lead to inappropriate node selections due to the lack of visualization-specific knowledge. To alleviate these issues, we introduce a rule-based pruning algorithm.

The key aspect of this algorithm is integrating domain knowledge, such as *data transformation rules* and *visualization rules*, to ensure the generated visualizations are syntactically correct. For example, after completing a GROUP BY operation, if an aggregate

function needs to be chosen for the y -axis, and the y -axis field is categorical, then functions such as SUM, AVG, and NONE become inapplicable. Our algorithm aims to exclude these inappropriate options to enhance the efficiency and accuracy of the search.

Furthermore, we present a function $L(S, A)$ that takes the current state S and the candidate operation set A as inputs, and outputs a set of operations meeting specific constraints, formalizing the pruning process based on domain-specific knowledge for visualizations.

Adaptive Random Exploration Strategy. During the visualization generation process, each selected visualization operation (e.g., chart type) significantly affects the final visualization. Although the UCB algorithm is designed to make the best choice based on the current state, it may tend to over-explore high-scoring branches as the number of simulations accumulates, overlooking other possibilities. Moreover, the UCB algorithm might favor sub-optimal or inefficient visualization actions, especially when the initial attempts are inaccurately evaluated by the reward function [12, 51].

To address these issues, we propose an adaptive random exploration strategy. This strategy adjusts the random selection probability based on the number of already selected clauses, facilitating the transition from exploitation to exploration. This approach increases search space coverage and diversity while avoiding local optima. It mirrors user behavior in constructing visualization queries, with extensive exploration at the initial stage and more targeted selection based on earlier trials at later stages.

The optimized formula for the selection strategy is as follows:

$$a_i = \begin{cases} \text{random}(L(S_i, A_i)), & P_{\text{exploration}} \\ \arg \max_{b \in L(S_i, A_i)} \left[Q_i(b) + c \sqrt{\frac{2 \ln t}{N_i(b)}} \right], & P_{\text{exploitation}} \end{cases} \quad (2)$$

In our proposed decision-making strategy, the final choice a_i for each node i is based on a series of parameters and functions. The set of legal clauses available for node i is represented by $L(S_i, A_i)$. The average reward for each clause b is estimated by $Q_i(b)$, while $N_i(b)$ tracks the number of times clause b has been visited. The total number of visits to the current node is denoted by t , and the constant c is used to balance exploration in the process.

Furthermore, the balance between exploration and exploitation in the strategy is controlled by two main probability parameters: $P_{\text{exploration}}$ and $P_{\text{exploitation}}$. The exploration probability, $P_{\text{exploration}}$, is defined as $p_n \alpha^n$, where p_n is the initial probability of making a random choice and α is a constant less than 1 that controls the decay rate of the random selection probability. On the other hand, the exploitation probability, $P_{\text{exploitation}}$, is defined as $1 - p_n \alpha^n$. As the number of clauses n selected in node i increases, the probability of random selection gradually decreases, while the probability of selection based on the UCB algorithm correspondingly increases.

5 LEARNING-TO-RATE VISUALIZATION

Unlike environments for games like Atari [45] or Go [50], which have clear reward and punishment rules, visualization evaluation lacks such clarity and can be biased if based on a single criterion.

To address this, we propose a composite reward function that incorporates visualization best practices, data features, and user preferences to comprehensively evaluate visualization quality.

5.1 Composite Reward Function

The composite reward function incorporates *rules of thumb* for visualizations, data features, and user preferences to evaluate the quality of visualization results. The formula for the Composite Reward Function (CRF) is as follows:

$$CRF = S_k \times (\beta S_d + (1 - \beta) S_u) \quad (3)$$

where S_k represents the assessment result from visualization domain knowledge, S_d indicates the score based on data features, and S_u denotes the user preference score. When evaluating a visualization result, we first conduct a preliminary evaluation based on visualization domain knowledge. If the result aligns with the domain knowledge, S_k is assigned a value of 1; otherwise, it is 0. Clearly, when $S_k = 0$ (i.e., the visualization result does not align with domain knowledge), the reward value CRF becomes 0. Conversely, when $S_k = 1$, the calculation of CRF depends on the weighted combination of S_d and S_u , where the weight coefficient β is introduced to balance the importance between data characteristic scoring and user preference scoring, which is set based on experience.

5.2 Learn Composite Reward Function

We leverage well-established, human-annotated visualization corpora [19, 39] and domain knowledge from the visualization community to learn the composite reward function.

Leveraging Domain Knowledge. We employ a rule-based method to ensure that the generated visualizations align with best practices and accurately reflect data features. Specifically, this method leverages the data selection, transformation, and visualization rules used by DeepEye [31].

For example, the rules for pie charts include: (1) Data selection rules: the x -axis should represent categorical data, while the y -axis should represent numerical data. (2) Data transformation rules: pie charts are unsuitable for data aggregated with the AVG operation, as they are primarily used to show proportions. (3) Visualization rules: the y -axis values should not include negatives, and there should be at least two distinct x -axis values to convey meaningful information. Using this approach, we design 15 rules for different types of visualizations and use these rules as strict constraints to filter out *low-quality* visualizations.

Capturing Data Features. Our model is trained on a dataset annotated by real users [31], containing 285,236 visualizations and their scores from 42 different domains. For each visualization, we analyze key features, including the data types of the x - and y -axes, the number of rows, extremes (maximum and minimum values), value diversity (number of different values and the ratio of unique values), and the correlation between the data of the x - and y -axes and the type of chart. We extract 14 core features for model training.

We use LambdaMART [67] to train our scoring model, which maps complex relationships between features and scores by building decision trees. Each tree evaluates features, such as the type of data on the x -axis, then examines other features, and finally outputs a score. This method enables the model to understand the deep connections between feature combinations and scores. When a new visualization is input, the model computes its features and uses the trained LambdaMART to predict a score.

Learning Common User Preferences. Capturing user preferences for visualization is essential in visualization recommendation systems. Although it is ideal to consider individual preferences in online recommendations, this poses a challenge due to the cold start problem – new users often lack sufficient interaction history, especially in visualization scenarios. In addition, collecting real-time interaction data and fine-tuning models online for each user can be resource-intensive and time-consuming.

To address these challenges, *our goal is to first learn common user preferences offline* and then make online adjustments to the MCGS algorithm based on feedback from visualization hints, which we will discuss in Section 6. This approach strikes a balance between recommending high-quality visualizations and fine-tuning the recommendations to better suit individual users.

To achieve this, our system utilizes the Generative Adversarial Network (GAN) [17], specifically the IRecGAN [5], to learn common user preferences using a real-world visualization corpus. This corpus includes user interaction logs collected from the Plotly community [39]. The dataset comprises historical interaction records of 17,469 users, with an average of 5.41 datasets and 1.85 visualizations per user, and each dataset contains an average of 24.39 attributes. For effective learning of user preferences across datasets and visualizations, we extracted 606 visualization configurations from the original dataset, inspired by [39]. Each configuration includes dataset-independent design choices such as chart types, colors, and sizes, enabling the model to learn preferences without direct dataset associations and to identify patterns from design choices. By aggregating data, user, and visualization features across different datasets, we generated 15,531 training samples to train the IRecGAN model [5].

In the first round of visualization recommendations, the evaluations obtained from the trained IRecGAN model serve as a crucial indicator of general user preferences. In the interaction phase (e.g., the second round), HAICHART refines and optimizes the recommendation strategy based on the visualization hints selected by users, which will be detailed in Section 6.

6 VISUALIZATION HINTS SELECTION

While the MCGS-based visualization generation algorithm can recommend high-quality visualizations, it might not always align with individual preferences. Conversely, interactive tools like Voyager2 [63] offer flexible exploration options but can overwhelm users with their complexity. These tools require users to sift through dense control panels and numerous visualization options, much like tackling a challenging “fill-in-the-blank” puzzle.

To address these issues, we introduce a visualization hints module. Visualization hints represent the high-level visualization intent derived by our system, as shown in Figure 7. This module simplifies the complex decision-making process of visualization creation into straightforward, user-friendly selections, akin to “multiple-choice questions”. This method collects user intent, allowing our MCGS algorithm to be fine-tuned online to better suit individual users.

Definition 6.1: [Visualization Hint] A visualization hint h corresponds to a visualization operation or action, such as selecting data fields, applying aggregate operations, and choosing chart types. It is expressed in easily understandable natural language. Each hint h

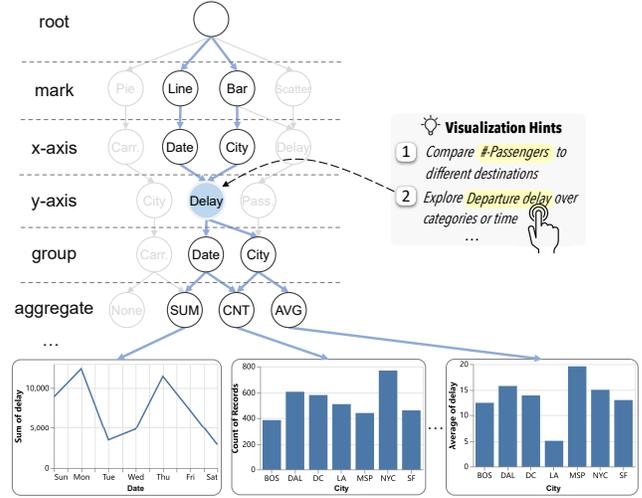


Figure 7: Visualization refinement with visualization hints

is associated with a set of visualizations $\mathbb{V} = \{v_1, v_2, \dots, v_n\}$, which are ranked based on the values of the composite reward function to prioritize high-value visualizations. \square

As shown in Figure 7, consider a data field-based hint: “Explore (data field name) over categories or time”, this hint essentially defines a partial visualization query relevant to the data field. Therefore, the user can browse these hints and select a desired hint for visualization. Next, HAICHART will generate a set of relevant visualizations based on the selected hint through the MCGS process. This approach not only simplifies the steps for the user to define a visualization query but also speeds up the process of gaining data insights by guiding the user through meaningful hints to quickly explore and understand the data.

Our goal is to select the top- k visualization hints that not only cover different aspects and ensure high-quality visualizations but also maintain an appropriate number of visualizations.

Definition 6.2: [Top- k Visualization Hints Selection] Given a set of hints $\mathbb{H} = \{h_1, h_2, \dots, h_n\}$, where each hint h_i is associated with a set of visualizations \mathbb{V}_i , and each visualization $v \in \mathbb{V}_i$ has an associated reward value r_v , the goal is to select a subset $\mathbb{H}' \subseteq \mathbb{H}$ consisting of k hints. The selection must maximize the total reward value of the subset \mathbb{H}' , while ensuring the total count of visualizations in \mathbb{H}' does not exceed a predefined budget B . The optimization problem can be formulated as follows:

$$\text{Maximize } F(\mathbb{H}') = \sum_{h_i \in \mathbb{H}'} \sum_{v \in \mathbb{V}_i} r_v \quad (4)$$

$$\text{Subject to } \sum_{h_i \in \mathbb{H}'} |\mathbb{V}_i| \leq B \quad \text{and} \quad |\mathbb{H}'| = k \quad (5)$$

where $|\mathbb{V}_i|$ represents the number of visualizations associated with hint h_i , $F(\mathbb{H}')$ is the total reward value of the selected subset \mathbb{H}' , $|\mathbb{H}'|$ is the number of selected hints, and B is the upper limit on the budget for the number of visualizations. \square

However, the problem of selecting hints is NP-hard because it is equivalent to a known NP-hard problem – the Budgeted Maximum

Algorithm 2: Top- k Visualization Hints Selection

Input: Set of hints $\mathbb{H} = \{h_1, h_2, \dots, h_n\}, B, k$;
Output: Selected top- k set of hints \mathbb{H}' ;

```
1  $\mathbb{H}' \leftarrow \emptyset$ ;  $totalCost \leftarrow 0$ ;  
2  $\mathbb{H}_v \leftarrow \{h \in \mathbb{H} \mid |h| \leq B\}$ ; // 1. Filter valid hints  
3  $\mathbb{H}_v \leftarrow \text{SortByScore}(\mathbb{H}_v)$ ; // 2. Sort hints by score  
4 // 3. Selection of top- $k$  hints  
5 for each  $h_i$  in  $\mathbb{H}_v$  do  
6   if  $|\mathbb{H}'| < k$  and  $totalCost + |h_i| \leq B$  then  
7      $\mathbb{H}'.append(h_i)$ ;  
8      $totalCost \leftarrow totalCost + |h_i|$ ;  
9   if  $|\mathbb{H}'| = k$  then  
10    break;  
11 return  $\mathbb{H}'$ ;
```

Coverage problem [20]. To address this, we propose an efficient algorithm to select the top- k visualization hints. To achieve effective hint selection, the first step is to construct a comprehensive set of hints and evaluate the benefit of each hint. The purpose of this initial phase is to ensure that in the subsequent selection process, the best hint can be selected from this evaluated set, thereby maximizing the reward value. In this way, the algorithm mainly consists of two steps: candidate hints generation and top- k hints selection.

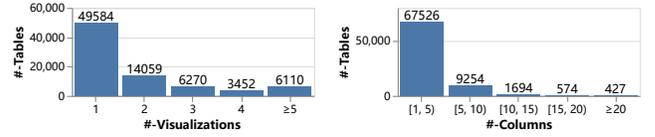
Candidate Visualization Hints Generation. We identify high-value nodes (*i.e.*, visualization operations) from our visualization query graph to pick the candidate hints. When computing the reward for each hint, we introduce a decay coefficient δ to adjust the scoring weight of each visualization, considering that the same visualization appearing in different hints may reduce its uniqueness. The decay coefficient δ , is defined as: $\delta = \log\left(\frac{N_{total}}{N_{viz}}\right)$, where N_{total} is the total number of hints and N_{viz} is the number of hints containing the specific visualization. This coefficient reflects the frequency of visualizations, with higher decay coefficients for more frequent visualizations, thereby reducing their overall scores. After multiplying this score by the reward value of the visualization, its final value is determined. Therefore, after collecting all the hints, we calculate the frequency of the appearance of each visualization and combine it with the decay coefficients to compute its score.

Top- k Visualization Hints Selection. Our goal is to select the best subset of hints that maximizes the total reward within a predefined budget B . Algorithm 2 shows the pseudo-code. It first selects all hints with a cost not exceeding the budget B , forming a candidate hint set (Line 2). Then, it sorts valid hints based on the corresponding average visualization score computed by the composite reward function (Line 3). The algorithm continues to go through this sorted set, picking hints to add to the final selected set until one of two conditions is met: the number of selected hints reaches the predefined k , or adding more hints would cause the total cost to exceed the budget B (Line 5-10). In this way, the algorithm prioritizes high-scoring hints while maintaining the budget constraint, thus maximizing the total reward value under the given budget.

User Feedback for Refinement. HAICHART leverages user-selected visualization hints to guide the node exploration strategy during the MCGS process. When users select a specific hint, they

Table 2: Statistics of the experimental datasets (Vis.: Charts)

Datasets	#-Tables	#-Vis.	Avg(#-Vis.)	Avg(#-Rows)	Avg(#-Col.)	Max(#-Col.)
VizML	79,475	162,905	2	2,817.8	3.3	25
KaggleBench	8	252	31.5	32,585.9	9.1	15

**Figure 8: Statistics of the VizML dataset****Table 3: Details in KaggleBench dataset**

	Datasets	#-Rows	#-Columns	#-Vis.
D1	StudentPerformance	1,000	8	34
D2	AirplaneCrashes	5,191	6	22
D3	VideoGame	6,825	15	48
D4	GooglePlayStore	9,659	11	33
D5	AvocadosPrice	18,249	7	20
D6	SuicideRates	27,820	9	25
D7	Zomato	51,717	11	49
D8	GunViolence	140,226	6	21

direct the search graph to expand in a certain direction. For example, as shown in Figure 7, if a user clicks on the hint “Explore delay over categories or time”, the system focuses on applying the data field “delay” to the y -axis. During the search process, the system freezes other nodes and explores only those related to “delay”, ensuring the generated results are relevant to the target field. Thus, this method can align the search process with user preferences and effectively prune the search space, enhancing the efficiency and accuracy of MCGS-based visualization recommendations.

7 EXPERIMENTS

7.1 Experiment Settings

Datasets. Table 2 shows two real-world datasets for experiments.

(1) VizML [19], derived from the Plotly community, features around 120,000 dataset-visualization pairs created by real users. This dataset was refined by removing entries missing table or chart data and cleaning up invalid characters. The dataset has four types of charts, namely bar, pie, line, and scatter visualizations, resulting in 79,475 valid dataset-visualization pairs. These were randomly divided into training, validation, and testing sets in a 7:1:2 ratio, allocating 55,632 pairs for training, 7,947 for validation, and 15,896 for testing. The statistical information of the dataset, shown in Figure 8, indicates a broad coverage, including various tables and visualizations used by users across different domains and tasks.

(2) KaggleBench [16] is a public benchmark designed to evaluate the effectiveness of visualization recommendations. Its data mainly come from numerous data competitions and the corresponding visualizations provided by the Kaggle platform. This dataset was refined by filtering out low-quality datasets, removing rows with

Table 4: Effectiveness of the first round of visualization recommendation

D	Tasks	Metrics	The State-of-the-Art Methods					Our Methods			
			Data2Vis [15]	VizGRank [16]	DeepEye [31]	PVisRec [39]	VizML [19]	LLM4Vis [59]	MCTS	HAICHART-	HAICHART
VizML	Data Queries	Hit@1	47.5%	57.6%	52.4%	52.3%	-	-	78.3%	79.7%	79.3%
		Hit@3	51.3%	67.2%	67.6%	58.7%	-	-	88.2%	91.3%	91.9%
	Design Choices	Hit@1	41.7%	34.9%	34.1%	28.9%	28.7%	47.9%	42.4%	50.6%	48.7%
		Hit@3	43.7%	42.9%	40.7%	51.3%	-	-	77.1%	81.8%	81.5%
	Overall	Hit@1	24.3%	25.6%	25.7%	21.8%	-	-	33.1%	37.9%	36.9%
		Hit@3	26.9%	30.1%	33.9%	42.3%	-	-	64.7%	68.4%	67.4%
KaggleBench	Data Queries	P@10	41.2%	58.7%	62.5%	42.5%	-	-	52.2%	60.0%	63.8%
		R10@30	25.0%	50.0%	48.7%	67.5%	-	-	73.6%	80.1%	83.7%
	Design Choices	P@10	88.7%	87.5%	93.7%	91.9%	Hit@2:78.3%	Hit@2:87.6%	93.8%	96.3%	96.3%
		R10@30	95.0%	81.3%	95.0%	85.0%	-	-	92.5%	96.2%	96.2%
	Overall	P@10	28.7%	43.7%	48.7%	36.7%	-	-	45.4%	51.3%	55.0%
		R10@30	13.8%	41.3%	33.7%	60.0%	-	-	63.8%	72.5%	74.9%

missing data, and fixing invalid characters. Finally, we obtained 8 datasets for evaluation, as shown in Table 3.

We trained or configured all methods using the VizML dataset and tested them on the two datasets mentioned above.

Methods. We evaluate the following methods.

(I) AI-powered Visualization Methods:

(1) Data2Vis [15] transforms datasets into visualization queries using a sequence-to-sequence model.

(2) VizGRank [16] achieves visualization recommendation by modeling the relationships between visualizations as graphs and using graph-based ranking algorithms.

(3) DeepEye [31] combines data features with domain knowledge to recommend top- k good visualizations.

(4) PVisRec [39] recommends a set of visualizations by learning from user preferences.

(5) VizML [19] uses deep learning models for visualization recommendations. It focuses on five specific types of tasks, excluding data querying. Our evaluation will assess VizML’s effectiveness in recommending design choices.

(6) LLM4Vis [59] uses in-context learning to interact with ChatGPT for recommending visualizations. Like VizML, it excludes data querying tasks. Thus, our comparison focuses on design choices, evaluated using the same Hit@ k metric as LLM4Vis.

(II) Human-powered Visualization Methods:

(7) Voyager2 [63] is an interactive system that allows users to create and explore visualizations through click-based interactions.

(III) Human and AI paired Visualization Methods:

(8) HAICHART (ours) is the full implementation based on MCGS and the composite reward function, as described in this paper.

(9) HAICHART- (ours) differs from HAICHART in the composite reward function, where the capturing data features model is trained on the VizML corpus.

(10) LLM4Vis+ (ours) is an improved version of LLM4Vis [59]. We enhanced LLM4Vis to support data queries and top- k visualizations.

(11) MCTS-based Baseline (ours) follows the same implementation as HAICHART, except it uses the MCTS-based method for recommending visualizations.

Metrics. Following previous studies evaluating automatic visualization systems [15, 19, 37, 69], we employ Hit@ k , P@ k , and Rt@ k as evaluation metrics. Given that creating a visualization involves

data queries, design choices, and final integration, our evaluation is divided into three tasks. The metrics are defined as follows:

(1) *Hit@ k* : This metric evaluates whether the ground truth appears in the top- k results. We apply this to the VizML dataset with k set to 3, as each user typically creates about 2 visualizations on average.

(2) *P@ k* : This metric measures how many ground truths are in the top- k results. For the KaggleBench dataset, we set k to 10, considering that each dataset typically contains over 20 visualizations.

(3) *Rt@ k* : This metric assesses how many of the top- t ground truths are covered in the top- k results. We use R10@30 for the KaggleBench dataset, analyzing how many of the top 10 ground truths are included in the first 30 results returned by the model.

Experimental Environment. Experiments were performed on an Ubuntu 22.04 Server LTS with dual Intel Xeon 8383C CPUs, 512GB RAM, and eight NVIDIA RTX 4090 GPUs.

7.2 Experimental Results

7.2.1 Exp-1: Effectiveness of the First-round of Recommendations. This experiment evaluates the effectiveness of the first-round visualization recommendations by HAICHART. We tested all methods on the VizML and KaggleBench datasets. Table 4 shows the results.

(1) Overall, our methods (HAICHART and HAICHART-) significantly outperform all state-of-the-art methods across all metrics, showing the effectiveness of our framework. The performance gap between HAICHART- and HAICHART is primarily due to the capturing data features module: HAICHART- is trained on the VizML training set, while HAICHART is trained on a more diversified corpus [31].

(2) On the VizML dataset, HAICHART achieves 36.9% Hit@1 and 67.4% Hit@3 on the overall task, surpassing the competitive methods DeepEye and PVisRec by 11.2% and 25.1%, respectively. On the KaggleBench dataset, HAICHART achieves 55% P@10 and 74.9% R10@30 on the overall task, outperforming DeepEye and PVisRec by 6.3% and 14.9%. These results demonstrate HAICHART’s effectiveness in leveraging MCGS algorithms with a composite reward function to find high-quality visualizations in a large search space.

(3) DeepEye performs well on the P@10 metric using a rule-based approach, while PVisRec excels on the Rt@30 metric with personalized recommendations. By blending visualization rules, user

Table 5: Performance vs. #-Iterations on KaggleBench

Tasks	Metrics	DeepEye	Voyager2			LLM4Vis+ (ours)			HAICHART (ours)		
		Iter. 1	Iter. 1	Iter. 2	Iter. 3	Iter. 1	Iter. 2	Iter. 3	Iter. 1	Iter. 2	Iter. 3
Data Queries	P@10	62.5%	45.0%	55.1%	58.0%	47.2%	65.0%	74.3%	63.8%	69.5%	79.2%
Design Choices	P@10	93.7%	78.7%	96.3%	97.4%	75.6%	91.5%	97.2%	96.3%	97.6%	99.3%
Overall	P@10	48.7%	40.0%	44.9%	45.7%	41.0%	55.6%	65.3%	55.0%	58.2%	68.8%

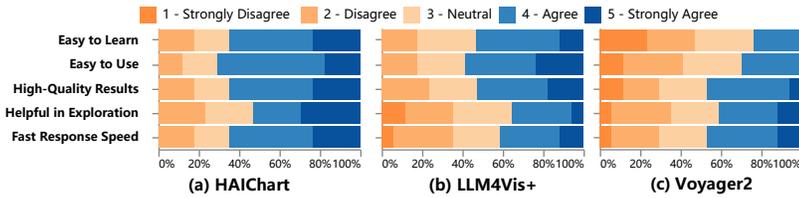


Figure 9: Qualitative analysis on user study

preferences, and data features, HAICHART achieves strong results in both the P@10 and Rt@30 metrics.

7.2.2 Exp-2: Effectiveness of Multi-round Recommendations. This experiment assesses whether HAICHART improves visualization recommendations and enhances user efficiency in data exploration through multiple rounds of interaction and visualization hints.

Participants. The study involved 17 participants (7 females, 10 males, aged 21-33), including 12 experts (6 Ph.D. candidates, 4 master’s students, and 2 undergraduates in computer science) and 5 non-experts from non-technical backgrounds.

Task. Participants used HAICHART, DeepEye, Voyager2, and LLM4Vis+ to analyze eight datasets from KaggleBench. These datasets required specific analytical tasks, such as analyzing the relationship between student performance and background factors in the StudentsPerformance dataset, as well as open-ended explorations to evaluate each tool’s adaptability and efficiency.

Procedure. (1) Preparation: Participants were introduced to the datasets and tools through a hands-on demonstration. (2) Experimentation: Participants explored datasets using these tools, during which we recorded their time spent, interactions, and visualizations. After each round, HAICHART offered the top-9 hints and new recommendations. (3) User Feedback: After the experiment, participants rated each tool on a five-point Likert scale [3, 25] and provided feedback through interviews.

User Study Results. We conducted both quantitative and qualitative analyses. The observations are as follows:

(1) As shown in Table 5, HAICHART improves visualization recommendations through user feedback. After two rounds of interactions, P@10 increased from 55% to 68.8%, outperforming LLM4Vis+ and Voyager2 by 3.5% and 23.1%, respectively. Unlike Voyager2’s manual exploration and LLM4Vis+’s time-consuming natural language queries, HAICHART’s hints guide users efficiently. DeepEye, achieving 48.7% in the first round, lacks multi-turn recommendation support and does not improve with additional interactions.

(2) According to Figure 10, HAICHART is 1.8 and 2.2 times faster than LLM4Vis+ and Voyager2, respectively, achieving an average recall of 83.7%. This is 21% and 24.8% higher than LLM4Vis+ and Voyager2, indicating that HAICHART meets analytical needs with fewer

Table 6: Effectiveness of Hints Selection

Dataset	Metrics	Round 1	Round 2	Round 3
KaggleBench	Hit@1	64.7%	65.2%	69.1%
	Hit@3	79.4%	82.2%	85.7%
	Hit@5	88.2%	89.6%	92.1%

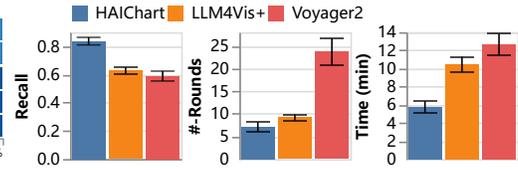


Figure 10: Quantitative analysis on user study

interactions and less time. Furthermore, user feedback (Figure 9) underscores HAICHART’s ease of learning and usability, facilitating more effective exploration and higher quality visualizations.

(3) Feedback from participants highlighted HAICHART’s practicality in data exploration. They widely acknowledged that HAICHART simplifies tasks and enhances efficiency with helpful hints. For example, Participant P1 stated, “Unlike Voyager2, which is not user-friendly for beginners, HAICHART makes it easier to choose data fields and operations through intuitive hints.” Participant P5 mentioned a limitation with LLM4Vis+, “Natural language interaction simplifies creating initial visualizations, but fine-tuning is challenging. LLMs often misunderstand my intent, requiring repeated adjustments.” This suggests that natural language ambiguity may cause ‘communication’ issues with LLMs during data exploration. Intuitive control panels or guidance, like hints, would enhance the user experience.

Nevertheless, participants also suggested improvements for HAICHART, such as adding a control panel for enhanced flexibility (Participant P10) and supporting more complex chart types to broaden analysis capabilities (Participant P15).

7.2.3 Exp-3: Effectiveness of Visualization Hints Selection. This experiment evaluates the effectiveness of visualization hint selection by analyzing interaction logs from 17 users in a multi-round recommendation experiment (Exp-2), using the Hit@k metric.

The results in Table 6 show that Hit@1 is 64.7% in the first round of recommendations, demonstrating the system’s ability to accurately predict the hint most interesting to the user. As the number of interaction rounds increases, the recommendation accuracy improves further. By the third round, the Hit@3 value increased to 85.7%. These results demonstrate the system’s effectiveness in prioritizing the information required by the user and the effectiveness of the hint selection in guiding the user’s decision-making process.

7.2.4 Exp-4: Ablation Study of the MCGS Optimization Techniques. We proposed several MCGS optimization techniques, including Rule-based Pruning and an Adaptive Random Exploration Strategy. We conducted ablation studies using the VizML and KaggleBench datasets to evaluate their impact on overall performance.

The results in Table 7 show that removing any optimization technique from either dataset leads to a decline in performance, confirming their importance. Specifically, in the VizML dataset,

Table 7: Ablation studies on HAICHART (overall performance)

Methods		VizML		KaggleBench	
		Hit@1	Hit@3	P@10	R10@30
HAICHART		36.9%	67.4%	55.0%	74.9%
Opt. Tech.	w/o Rule-based Pruning	34.6%	65.3%	37.4%	40.0%
	w/o Adapt. Random Exploration	33.4%	42.1%	45.7%	65.0%
Composite Reward Func.	w/o Domain Knowl.	26.3%	60.2%	31.0%	33.8%
	w/o User Preferences	30.8%	64.2%	40.7%	54.9%
	w/o Data Features	34.2%	64.1%	36.0%	68.7%

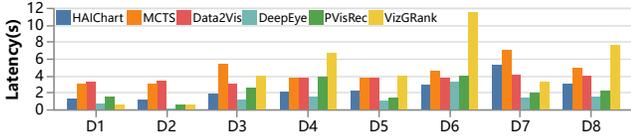


Figure 11: Efficiency on KaggleBench datasets

removing the Adaptive Random Exploration Strategy reduces the improvement from Hit@1 to Hit@3 from 30.5% to 8.7%, which demonstrates that this strategy enhances result diversity and avoids local optima, thereby improving performance.

7.2.5 Exp-5: Ablation Study of the Composite Reward Function. Our composite reward function consists of three components: data features, visualization domain knowledge, and user preferences. We conducted ablation studies using the VizML and KaggleBench datasets to evaluate their impact on overall performance.

Table 7 shows the results with the following observations.

Overall, removing any component decreased the performance of HAICHART. Performance drops were more significant when removing domain knowledge or user preferences, indicating that these aspects are crucial for aligning visualizations with user needs. In summary, integrating data features, visualization domain knowledge, and user preferences in the composite reward function enhances the effectiveness of visualization recommendations.

7.2.6 Exp-6: The Efficiency of HAICHART. We compared the efficiency of HAICHART and existing methods using the KaggleBench dataset, which includes eight tables (D1 to D8) of varying sizes.

Figure 11 shows the results with the following observations:

- (1) HAICHART generates visualizations in 1 to 5 seconds, with an average of 2.4 seconds, and hint generation takes only 1 to 3 milliseconds, validating its capability for rapid visualization exploration.
- (2) HAICHART is, on average, 1.8 times faster than MCTS, especially for datasets with many columns. For example, with the D3 dataset (15 columns), HAICHART’s graph-based structure outperforms MCTS’s tree-based structure by 2.9 times.

Overall, HAICHART is efficient enough for data visualization.

8 RELATED WORK

Human-powered Visualization tools such as Tableau, Voyager, and Polaris [1, 2, 44, 52, 62, 66] allow users to select or adjust data sources, chart types, and data transformation operations for data visualization. For example, Tableau [2] offers a code-free interface for creating visualizations through click-and-drag actions. However, these tools are highly dependent on user skills, presenting challenges such as steep learning curves.

AI-powered Automatic Visualization tools [8, 16, 19, 26, 28, 30–32, 39, 40, 69] uses algorithms to automatically generate and recommend meaningful visualizations. For example, DeepEye [31] uses machine learning to recommend good visualizations based on data features and domain knowledge. Similarly, VizML [19] trains a neural network model to recommend charts based on large-scale real-world visualization cases. PVisRec [39] suggests personalized visualizations based on the user’s past interactions. However, these systems may overlook user intent, resulting in recommended visualizations that may not always meet users’ needs.

Large Language Models for Visualization (LLM4VIS). LLM4VIS leverages large language models to transform natural language queries into data visualizations [33–36, 53, 56, 59, 60]. For example, ChartGPT [56] employs LLMs to generate visualizations from natural language queries, while LLM4Vis [59] uses few-shot learning to suggest visualization types and explain them in text. These methods rely on users providing clear query descriptions and may struggle with accurately modifying visualizations through natural language [70]. Our method complements LLM4VIS by providing visualization hints and easy-to-adjust features, reducing user input and making the visualization creation process more efficient.

Reinforcement Learning is a learning approach based on trial-and-error, where an agent receives feedback from its environment [27]. It has been successfully used in various tasks such as item recommendation [58] and data preparation [7–9]. Recently, it has also been applied to visualization tasks [10, 13, 64, 69]. For example, PI2 [10, 55] uses MCTS to generate interactive UI widgets from SQL logs, assisting developers in understanding necessary queries for analysis tasks. Unlike PI2, which requires user-provided SQL queries, HAICHART automatically recommends visualizations for datasets and offers hints to guide data exploration. Both systems facilitate data exploration but differ in their application scenarios.

9 CONCLUSION

We introduce HAICHART, which pairs human insight with AI capabilities to enhance visualization quality *progressively* and *iteratively* through user feedback. HAICHART utilizes a Monte Carlo Graph Search-based visualization algorithm for automatically recommending high-quality visualizations. It is also equipped with a visualization hints mechanism to actively incorporate the user feedback and thus fine-tune the visualization generation algorithm iteratively. HAICHART has been validated for its effectiveness in creating high-quality visualizations that align with user preferences. Future research could explore integrating LLMs into our framework to generate and evaluate visualizations, potentially enhancing generalizability and robustness for diverse applications.

ACKNOWLEDGMENTS

This paper was supported by Guangdong Basic and Applied Basic Research Foundation (2023A1515110545), National Key R&D Program of China (2023YFB4503600), NSF of China (61925205, 62232009, 62102215), Zhongguancun Lab, Huawei, TAL education, Beijing National Research Center for Information Science and Technology (BN-Rist), CCF-Huawei Populus Grove Fund (CCF-HuaweiDB202306), and the Fundamental Research Funds for the Central Universities.

REFERENCES

- [1] July 15, 2024. Excel. <https://www.microsoft.com/en-us/microsoft-365/excel>
- [2] July 15, 2024. Tableau. <https://www.tableau.com/>
- [3] I Elaine Allen and Christopher A Seaman. 2007. Likert scales and data analyses. *Quality progress* 40, 7 (2007), 64–65.
- [4] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47 (2002), 235–256.
- [5] Xueying Bai, Jian Guan, and Hongning Wang. 2019. A Model-Based Reinforcement Learning with Adversarial Training for Online Recommendation. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*. 10734–10745.
- [6] Richard Bellman. 1957. A Markovian decision process. *Journal of mathematics and mechanics* (1957), 679–684.
- [7] Chengliang Chai, Jiabin Liu, Nan Tang, Guoliang Li, and Yuyu Luo. 2022. Selective Data Acquisition in the Wild for Model Charging. *Proc. VLDB Endow.* 15, 7 (2022), 1466–1478.
- [8] Chengliang Chai, Nan Tang, Ju Fan, and Yuyu Luo. 2023. Demystifying Artificial Intelligence for Data Preparation. In *SIGMOD Conference Companion*. ACM, 13–20.
- [9] Chengliang Chai, Jiayi Wang, Yuyu Luo, Zeping Niu, and Guoliang Li. 2023. Data Management for Machine Learning: A Survey. *IEEE Trans. Knowl. Data Eng.* 35, 5 (2023), 4646–4667.
- [10] Yiru Chen and Eugene Wu. 2022. PI2: End-to-end Interactive Visualization Interface Generation from Queries. In *SIGMOD '22: International Conference on Management of Data, Philadelphia, PA, USA, June 12 - 17, 2022*. ACM, 1711–1725.
- [11] Liying Cheng, Xingxuan Li, and Lidong Bing. 2023. Is GPT-4 a Good Data Analyst? *arXiv preprint arXiv:2305.15038* (2023).
- [12] Johannes Czech, Patrick Korus, and Kristian Kersting. 2021. Improving AlphaZero Using Monte-Carlo Graph Search. In *Proceedings of the Thirty-First International Conference on Automated Planning and Scheduling, ICAPS 2021, Guangzhou, China (virtual), August 2-13, 2021*. AAAI Press, 103–111.
- [13] Dazhen Deng, Aoyu Wu, Huamin Qu, and Yingcai Wu. 2022. DASHBOT: Insight-driven dashboard generation based on deep reinforcement learning. *IEEE Transactions on Visualization and Computer Graphics* 29, 1 (2022), 690–700.
- [14] Victor Dibia. 2023. LIDA: A Tool for Automatic Generation of Grammar-Agnostic Visualizations and Infographics using Large Language Models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics: System Demonstrations, ACL 2023, Toronto, Canada, July 10-12, 2023*. Association for Computational Linguistics, 113–126.
- [15] Victor Dibia and Çağatay Demiralp. 2019. Data2vis: Automatic generation of data visualizations using sequence-to-sequence recurrent neural networks. *IEEE computer graphics and applications* 39, 5 (2019), 33–46.
- [16] Qianfeng Gao, Zhenying He, Yinan Jing, Kai Zhang, and X. Sean Wang. 2021. VizGRank: A Context-Aware Visualization Recommendation Method Based on Inherent Relations Between Visualizations. In *Database Systems for Advanced Applications - 26th International Conference, DASFAA 2021, Taipei, Taiwan, April 11-14, 2021, Proceedings, Part III (Lecture Notes in Computer Science, Vol. 12683)*. Springer, 244–261.
- [17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2020. Generative adversarial networks. *Commun. ACM* 63, 11 (2020), 139–144.
- [18] Yi He, Shixiong Cao, Yang Shi, Qing Chen, Ke Xu, and Nan Cao. 2024. Leveraging Large Models for Crafting Narrative Visualization: A Survey. *arXiv preprint arXiv:2401.14010* (2024).
- [19] Kevin Hu, Michiel A Bakker, Stephen Li, Tim Kraska, and César Hidalgo. 2019. Vizml: A machine learning approach to visualization recommendation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [20] Samir Khuller, Anna Moss, and Joseph Naor. 1999. The Budgeted Maximum Coverage Problem. *Inf. Process. Lett.* 70, 1 (1999), 39–45.
- [21] Nam Wook Kim, Grace Myers, and Benjamin Bach. 2023. How Good is ChatGPT in Giving Advice on Your Visualization Design? *arXiv preprint arXiv:2310.09617* (2023).
- [22] Doris Jung-Lin Lee, Dixin Tang, Kunal Agarwal, Thyne Boonmark, Caitlyn Chen, Jake Kang, Ujjaini Mukhopadhyay, Jerry Song, Micah Yong, Marti A Hearst, et al. 2021. Lux: always-on visualization recommendations for exploratory dataframe workflows. *arXiv preprint arXiv:2105.00121* (2021).
- [23] Edouard Leurent and Odalric-Ambrym Maillard. 2020. Monte-Carlo Graph Search: the Value of Merging Similar States. In *Proceedings of The 12th Asian Conference on Machine Learning, ACML 2020, 18-20 November 2020, Bangkok, Thailand (Proceedings of Machine Learning Research, Vol. 129)*. PMLR, 577–592.
- [24] Boyan Li, Yuyu Luo, Chengliang Chai, Guoliang Li, and Nan Tang. 2024. The Dawn of Natural Language to SQL: Are We Fully Ready? *Proc. VLDB Endow.* 17 (2024).
- [25] Guozheng Li, Runfei Li, Yunshan Feng, Yu Zhang, Yuyu Luo, and Chi Harold Liu. 2024. Colnsight: Visual Storytelling for Hierarchical Tables with Connected Insights. *IEEE Transactions on Visualization and Computer Graphics* (2024).
- [26] Haotian Li, Yong Wang, Songheng Zhang, Yangqiu Song, and Huamin Qu. 2021. KG4Vis: A knowledge graph-based approach for visualization recommendation. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2021), 195–205.
- [27] Jiabin Liu, Chengliang Chai, Yuyu Luo, Yin Lou, Jianhua Feng, and Nan Tang. 2022. Feature Augmentation with Reinforcement Learning. In *ICDE*. IEEE, 3360–3372.
- [28] Yuyu Luo, Chengliang Chai, Xuedi Qin, Nan Tang, and Guoliang Li. 2020. Interactive Cleaning for Progressive Visualization through Composite Questions. In *ICDE*. IEEE, 733–744.
- [29] Yuyu Luo, Chengliang Chai, Xuedi Qin, Nan Tang, and Guoliang Li. 2020. VisClean: Interactive Cleaning for Progressive Visualization. *Proc. VLDB Endow.* 13, 12 (2020), 2821–2824.
- [30] Yuyu Luo, Xuedi Qin, Chengliang Chai, Nan Tang, Guoliang Li, and Wenbo Li. 2022. Steerable Self-Driving Data Visualization. *IEEE Trans. Knowl. Data Eng.* 34, 1 (2022), 475–490.
- [31] Yuyu Luo, Xuedi Qin, Nan Tang, and Guoliang Li. 2018. DeepEye: Towards Automatic Data Visualization. In *34th IEEE International Conference on Data Engineering, ICDE 2018, Paris, France, April 16-19, 2018*. IEEE Computer Society, 101–112.
- [32] Yuyu Luo, Xuedi Qin, Nan Tang, Guoliang Li, and Xinran Wang. 2018. DeepEye: Creating Good Data Visualizations by Keyword Search. In *SIGMOD Conference*. ACM, 1733–1736.
- [33] Yuyu Luo, Jiawei Tang, and Guoliang Li. 2021. nvBench: A Large-Scale Synthesized Dataset for Cross-Domain Natural Language to Visualization Task. *NL2VIS, IEEE Visualization Conference Companion* abs/2112.12926 (2021).
- [34] Yuyu Luo, Jiawei Tang, Guoliang Li, and Chengliang Chai. 2021. Empowering natural language to visualization neural translation using synthesized benchmarks. *IEEE Visualization Conference* (2021).
- [35] Yuyu Luo, Nan Tang, Guoliang Li, Chengliang Chai, Wenbo Li, and Xuedi Qin. 2021. Synthesizing Natural Language to Visualization (NL2VIS) Benchmarks from NL2SQL Benchmarks. In *SIGMOD Conference*. ACM, 1235–1247.
- [36] Yuyu Luo, Nan Tang, Guoliang Li, Jiawei Tang, Chengliang Chai, and Xuedi Qin. 2022. Natural Language to Visualization by Neural Machine Translation. *IEEE Trans. Vis. Comput. Graph.* 28, 1 (2022), 217–226.
- [37] Yuyu Luo, Yihui Zhou, Nan Tang, Guoliang Li, Chengliang Chai, and Leixian Shen. 2023. Learned Data-aware Image Representations of Line Charts for Similarity Search. *Proc. ACM Manag. Data* 1, 1 (2023), 88:1–88:29.
- [38] Dominik Moritz, Chenglong Wang, Gregory Nelson, Halden Lin, Adam M. Smith, Bill Howe, and Jeffrey Heer. 2019. Formalizing Visualization Design Knowledge as Constraints: Actionable and Extensible Models in Draco. *IEEE Trans. Visualization & Comp. Graphics (Proc. InfoVis)* (2019).
- [39] Xin Qian, Ryan A Rossi, Fan Du, Sungchul Kim, Eunye Koh, Sana Malik, Tak Yeon Lee, and Nesreen K Ahmed. 2022. Personalized visualization recommendation. *ACM Transactions on the Web (TWEB)* 16, 3 (2022), 1–47.
- [40] Xuedi Qin, Yuyu Luo, Nan Tang, and Guoliang Li. 2018. DeepEye: An automatic big data visualization framework. *Big data mining and analytics* 1, 1 (2018), 75–82.
- [41] Xuedi Qin, Yuyu Luo, Nan Tang, and Guoliang Li. 2020. Making data visualization more efficient and effective: a survey. *VLDB J.* 29, 1 (2020), 93–117.
- [42] Christian P Robert, George Casella, Christian P Robert, and George Casella. 1999. Monte carlo integration. *Monte Carlo statistical methods* (1999), 71–138.
- [43] Abdallah Saffidine, Tristan Cazenave, and Jean Méhat. 2012. UCD: Upper Confidence bound for rooted Directed acyclic graphs. *Knowledge-Based Systems* 34 (2012), 26–33.
- [44] Arvind Satyanarayan and Jeffrey Heer. 2014. Lyra: An Interactive Visualization Design Environment. *Comput. Graph. Forum* 33, 3 (2014), 351–360.
- [45] Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, Timothy P. Lillicrap, and David Silver. 2020. Mastering Atari, Go, chess and shogi by planning with a learned model. *Nat.* 588, 7839 (2020), 604–609.
- [46] Leixian Shen, Enya Shen, Yuyu Luo, Xiaocong Yang, Xuming Hu, Xiongshuai Zhang, Zhiwei Tai, and Jianmin Wang. 2023. Towards Natural Language Interfaces for Data Visualization: A Survey. *IEEE Trans. Vis. Comput. Graph.* 29, 6 (2023), 3121–3144.
- [47] Leixian Shen, Enya Shen, Zhiwei Tai, Yun Wang, Yuyu Luo, and Jianmin Wang. 2022. GALVIS: Visualization Construction through Example-Powered Declarative Programming. In *CIKM*. ACM, 4975–4979.
- [48] Yelong Shen, Jianshu Chen, Po-Sen Huang, Yuqing Guo, and Jianfeng Gao. 2018. M-walk: Learning to walk over graphs using monte carlo tree search. *Advances in Neural Information Processing Systems* 31 (2018).
- [49] Danqing Shi, Xinyue Xu, Fuling Sun, Yang Shi, and Nan Cao. 2021. Calliope: Automatic Visual Data Story Generation from a Spreadsheet. *IEEE Trans. Vis. Comput. Graph.* 27, 2 (2021), 453–463.
- [50] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Vedavyas Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy P. Lillicrap, Madeleine Leach, Koray

- Kavukcuoglu, Thore Graepel, and Demis Hassabis. 2016. Mastering the game of Go with deep neural networks and tree search. *Nat.* 529, 7587 (2016), 484–489.
- [51] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy P. Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. 2017. Mastering the game of Go without human knowledge. *Nat.* 550, 7676 (2017), 354–359.
- [52] Chris Stolte, Diane Tang, and Pat Hanrahan. 2008. Polaris: a system for query, analysis, and visualization of multidimensional databases. *Commun. ACM* 51, 11 (2008), 75–84.
- [53] Jiawei Tang, Yuyu Luo, Mourad Ouzzani, Guoliang Li, and Hongyang Chen. 2022. Sevi: Speech-to-Visualization through Neural Machine Translation. In *SIGMOD Conference*. ACM, 2353–2356.
- [54] Nan Tang, Chenyu Yang, Ju Fan, Lei Cao, Yuyu Luo, and Alon Y. Halevy. 2024. VeriFAI: Verified Generative AI. In *CIDR*. www.cidrdb.org.
- [55] Jeffrey Tao, Yiru Chen, and Eugene Wu. 2022. Demonstration of PI2: Interactive Visualization Interface Generation for SQL Analysis in Notebook. In *SIGMOD Conference*. ACM, 2365–2368.
- [56] Yuan Tian, Weiwei Cui, Dazhen Deng, Xinjing Yi, Yurun Yang, Haidong Zhang, and Yingcai Wu. 2023. Chartgpt: Leveraging llms to generate charts from abstract natural language. *arXiv preprint arXiv:2311.01920* (2023).
- [57] Marko Tot, Michelangelo Conserva, Diego Perez Liebana, and Sam Devlin. 2022. Turning Zeroes into Non-Zeroes: Sample Efficient Exploration with Monte Carlo Graph Search. In *IEEE Conference on Games, CoG 2022, Beijing, China, August 21-24, 2022*. IEEE, 300–306.
- [58] Kai Wang, Zhene Zou, Qilin Deng, Jianrong Tao, Runze Wu, Changjie Fan, Liang Chen, and Peng Cui. 2021. Reinforcement Learning with a Disentangled Universal Value Function for Item Recommendation. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*. AAAI Press, 4427–4435.
- [59] Lei Wang, Songheng Zhang, Yun Wang, Ee-Peng Lim, and Yong Wang. 2023. LLM4Vis: Explainable Visualization Recommendation using ChatGPT. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: EMNLP 2023 - Industry Track, Singapore, December 6-10, 2023*. Association for Computational Linguistics, 675–692.
- [60] Xingbo Wang, Furui Cheng, Yong Wang, Ke Xu, Jiang Long, Hong Lu, and Huamin Qu. 2022. Interactive data analysis with next-step natural language query recommendation. *arXiv preprint arXiv:2201.04868* (2022).
- [61] Matthew O Ward, Georges Grinstein, and Daniel Keim. 2010. *Interactive data visualization: foundations, techniques, and applications*. CRC press.
- [62] Kanit Wongsuphasawat, Dominik Moritz, Anushka Anand, Jock Mackinlay, Bill Howe, and Jeffrey Heer. 2015. Voyager: Exploratory analysis via faceted browsing of visualization recommendations. *IEEE transactions on visualization and computer graphics* 22, 1 (2015), 649–658.
- [63] Kanit Wongsuphasawat, Zening Qu, Dominik Moritz, Riley Chang, Felix Ouk, Anushka Anand, Jock Mackinlay, Bill Howe, and Jeffrey Heer. 2017. Voyager 2: Augmenting visual analysis with partial view specifications. In *Proceedings of the 2017 chi conference on human factors in computing systems*. 2648–2659.
- [64] Aoyu Wu, Wai Tong, Tim Dwyer, Bongshin Lee, Petra Isenberg, and Huamin Qu. 2020. Mobilevisfixer: Tailoring web visualizations for mobile phones leveraging an explainable reinforcement learning framework. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2020), 464–474.
- [65] Aoyu Wu, Yun Wang, Mengyu Zhou, Xinyi He, Haidong Zhang, Huamin Qu, and Dongmei Zhang. 2021. MultiVision: Designing analytical dashboards with deep learning based recommendation. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2021), 162–172.
- [66] Eugene Wu, Fotis Psallidas, Zhengjie Miao, Haoci Zhang, and Laura Rettig. 2017. Combining Design and Performance in a Data Visualization Management System. In *CIDR*. www.cidrdb.org.
- [67] Qiang Wu, Chris JC Burges, Krysta M Svore, and Jianfeng Gao. 2008. *Ranking, boosting, and model adaptation*. Technical Report. Citeseer.
- [68] Yilin Ye, Jianing Hao, Yihan Hou, Zhan Wang, Shishi Xiao, Yuyu Luo, and Wei Zeng. 2024. Generative AI for visualization: State of the art and future directions. *Visual Informatics* 8, 2 (2024), 43–66.
- [69] Mengyu Zhou, Qingtao Li, Xinyi He, Yuejiang Li, Yibo Liu, Wei Ji, Shi Han, Yining Chen, Daxin Jiang, and Dongmei Zhang. 2021. Table2Charts: Recommending charts by learning shared table representations. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 2389–2399.
- [70] Yizhang Zhu, Shiyin Du, Boyan Li, Yuyu Luo, and Nan Tang. 2024. Are Large Language Models Good Statisticians? *CoRR* abs/2406.07815 (2024).